Convex Parameterizations and Fidelity Bounds for Nonlinear Identification and Reduced-Order Modelling

Mark M. Tobenkin Ian R. Manchester Alexandre Megretski

arXiv:1701.06652v1 [cs.SY] 23 Jan 2017

***Abstract*—Model instability and poor prediction of long-term behavior are common problems when modeling dynamical systems using nonlinear "black-box" techniques. Direct optimization of the long-term predictions, often called *simulation error minimization*, leads to optimization problems that are generally non-convex in the model parameters and suffer from multiple local minima. In this work we present methods which address these problems through convex optimization, based on Lagrangian relaxation, dissipation inequalities, contraction theory, and semidefinite programming. We demonstrate the proposed methods with a model order reduction task for electronic circuit design and the identification of a pneumatic actuator from experiment.**

I. Introduction

Building approximate models of dynamical systems from data is a ubiquitous task in the sciences and engineering. Reduced order modelling and system identification are two well-known cases, which differ mainly in the source of the data. In reduced order modelling the data come from simulation "snapshots" of a complex model[1], which is usually based on physical principles (e.g. [2], [3], [4]). The central problem is finding a model that simulates accurately despite unmodelled dynamics. In system identification the data come from experimental recordings [5], [6], [7], [8], and this brings the additional problems of measurement noise and exogenous disturbances. When some known structure can be incorporated into a model it is usually preferable to do so, but black-box modeling plays an important role when first-principles models are either weakly identifiable, too complicated for the eventual application, or simply unavailable.

The basic objective we consider is to obtain a dynamical system that accurately reproduces the observed relationship between some given input and output data, and can reliably generalize to other inputs. A complete solution using optimization methods requires a specification of a model class, a choice of objective function, and ultimately an algorithm for minimizing the objective function over the model class.

The model class we propose consists of discrete-time nonlinear state space models. Such models have previously been studied in the context of model reduction (e.g. [3]) and system identification (e.g. [9]). It includes all linear time-invariant (LTI) systems, systems defined by LTI internal dynamics with nonlinear input/output mappings (e.g. Wiener and Hammerstein models [10]), and includes nonlinear autoregressive models (e.g. [5], [11]) by choosing the state to be a truncated history of outputs.

An alternative is to model the current output as a static function of a finite history of inputs, e.g. Nonlinear Finite-Impulse-Response (NFIR) and Volterra models [5], [12]. However, such models are inefficient at representing inherently dynamic behaviour such as resonance, which is naturally expressed by models with internal feedback.

The inclusion of feedback in a model means that stability is no longer guaranteed, and in many applications an unstable model would be unacceptable. The ability of a model to "generalize" to inputs not in the training data depends critically on model stability. Ensuring stability of models with feedback is a significant problem and has been studied for many years for linear system identification [8]. It is known that stability of linear autoregressive models can be guaranteed if the input is generated by an autoregressive process and the data sequence is sufficiently long [13]. Stability of linear state space models generated by subspace identification was addressed by padding a shifted state matrix with zeros in [14], and by regularization in [15]. A linear matrix inequality (LMI) parametrization of stable linear models was given in [16].

The challenge of ensuring stability of nonlinear models with feedback is discussed in [7], [8], [11]. To the authors' knowledge, there are few prior published methods that guarantee stable models regardless of the source of the data. Worthy of note are the recent results of [17] giving small-gain and passivity-like stability guarantees for reduction of a high-order linear system in feedback with a convergent nonlinear system.

The objective function we focus on in this paper is simulation error, a.k.a. output error. This function is in general highly non-convex for models with feedback, and it is challenging to compute the global minimum [8]. Indeed, one can find poor local minima for the class of second-order linear models [18].

A computationally simpler alternative is to minimize the *equation error*, i.e. the one-step ahead prediction error, rather than the long-term simulation error. Nonlinear ARX and subspace identification, for example, use these criteria [5], [6]. In many cases, the minimization amounts to basic least squares and can be solved very efficiently. The disadvantage is that there is in general no connection between the equation error and the simulation error, since an unstable (or weakly stable) model might be found for which small equation errors accumulate under recursive simulation.

The main contributions of this paper are a class of nonlinear state-space models that is convex in its parameters and guarantees stability, and a family of convex upper bounds

M. Tobenkin and A. Megretski are with the Department of Electrical Engineering and Computer Science, Massachusetts Institute of Technology, Cambridge, MA, 02139. Email: {mmt, ameg}@mit.edu

I. Manchester is with the Australian Centre for Field Robotics and School of Aerospace, Mechanical and Mechatronic Engineering, University of Sydney, NSW 2006, Australia. Email: ian.manchester@sydney.edu.au.

This was supported by National Science Foundation Grant No. 0835947, the LANL Information Science and Technology Institute, and the Australian Research Council (DP130100551).

[1]We note that another class of model reduction algorithms work directly from the defining equations of the high-order model rather than simulation snapshots, e.g. balanced truncation and Hankel-optimal reduction [1].

for minimization of simulation error. The combination of these allows problems in reduced order modelling and identification to be approached using semidefinite programming. Our approach builds upon the ideas proposed in [19] for nonlinear autoregressive models, which was extended and applied to electronics model reduction in [20]. The model set and some fidelity bounds in this paper were reported without proof in [21]. The present paper significantly extends these results by introducing more accurate fidelity bounds based on Lagrangian relaxation, complete proofs of technical results, and new experimental results. We do not discuss the effects of measurement noise, but we note that a modified version of the proposed method gives consistent estimates when repeated measurements are available [22], [23], [24].

II. Problem Setup & Preliminaries

Spaces of vector signals, i.e. mappings $\{0, 1, ..., T\} \to \mathbb{R}^n$ are denoted l_T^n. For a symmetric matrix $A = A'$, $A \geq 0$ and $A > 0$ mean that A is positive semidefinite and positive-definite, respectively. Inequalities $A \geq B, A > B$ mean $A - B \geq 0, A - B > 0$, respectively. For symmetric positive-semidefinite Q the notation $|v|_Q^2$ is shorthand for $v'Qv$. Throughout the paper, signals with a tilde, e.g. $\tilde{y}(t)$, refer to recorded data, whereas signals without, e.g. $y(t)$ refer to the corresponding quantity generated by an identified model.

Assume that training data are given in the form of finite sequences of inputs and outputs: $\tilde{u} = \{\tilde{u}(0), \tilde{u}(1), ..., \tilde{u}(T)\} \in l_T^m$, $\tilde{y} = \{\tilde{y}(0), \tilde{y}(1), ..., \tilde{y}(T)\} \in l_T^p$. We use the notation $\tilde{z} := \{\tilde{u}, \tilde{y}\}$ to represent the "available data".

In this paper, a model class is a triple $\mathcal{M} = (a, g, \Theta)$ where $a : \mathbb{R}^n \times \mathbb{R}^m \times \mathbb{R}^q \to \mathbb{R}^n$ and $g : \mathbb{R}^n \times \mathbb{R}^m \times \mathbb{R}^q \to \mathbb{R}^p$ are continuously differentiable functions, and $\Theta \subset \mathbb{R}^q$ is a set of q-dimensional parameter vectors. A model class (a, g, Θ) is associated with state-space dynamical systems of the form

$$x(t+1) = a(x(t), u(t), \theta), \quad (1)$$

$$y(t) = g(x(t), u(t), \theta), \quad (2)$$

where $x \in \mathbb{R}^n, u \in \mathbb{R}^m, y \in \mathbb{R}^p$, $\theta \in \Theta$. It is desirable that identified models generalize to inputs not in the training data, it is appropriate to impose a strong notion of stability that which ensures predictable response to a wide variety of inputs independent of model initial conditions, such as the following:

Definition 1: A system (1), (2) is *incrementally ℓ^2 stable* if for every pair of initial conditions $x_1(0), x_2(0)$ and every input sequence $u(t), t = 1, 2, ...$, the states $x_1(t), x_2(t)$ and outputs $y_1(t) = g(x_1(t), u(t))$ and $y_2(t) = g(x_2(t), u(t))$ satisfy $\sum_{t=0}^{\infty} |x_1(t) - x_2(t)|^2 < \infty$, $\sum_{t=0}^{\infty} |y_1(t) - y_2(t)|^2 < \infty$.

In addition to a model set, one must define a measure of model fidelity. There are many possibilities, but perhaps the most natural and straightforward measure of quality is *simulation error*, also known as *output error*:

$$J_{se}(\theta, x, \tilde{z}) := \sum_{t=0}^{T} |y(t) - \tilde{y}(t)|^2 \quad (3)$$

which is to be minimized over parameters θ and states x for given data $\tilde{z}$, subject to dynamic constraints (1), (2).

The major difficulties we seek to address in this paper are:

(a) Obtaining a useful set of stable models. For all non-trivial parameterizations of functions $a(x(t), u(t), \theta)$, an explicit and tractable characterization of the set of all θ defining a stable model is unknown.

(b) Efficient global minimization of J_{se}. Even if the cost function $J(\theta, x, \tilde{z})$ is convex in θ, the dynamic constraints (1) generally define a nonconvex set: the simulation output is generated by the recursion

$$y(t) = g(a(a(...a(x(0), \tilde{u}(0), \theta), ..., \tilde{u}(t-1), \theta), \tilde{u}(t), \theta),$$

so even if a and g are linear in the parameter θ, J_{se} is a highly nonlinear function of the model

In this paper, we construct a family of computationally tractable approximations to the simulation error minimization problem. More specifically, we formulate model classes $\hat{\mathcal{M}} = (a, g, \hat{\Theta})$ and optimization problems of the form: $\min_{\theta \in \hat{\Theta}} \hat{J}(\theta, \tilde{z})$, with the following properties:

(a) $\hat{\mathcal{M}}$ is a finite-dimensional convex parametrization of models satisfying the following inclusions: $LTI \subseteq \hat{\mathcal{M}} \subseteq \mathcal{M}_{\ell 2}$ where LTI refers to the set of all stable linear time invariant models of appropriate dimension, and $\mathcal{M}_{\ell 2}$ is the set of all (possibly nonlinear) models that are *incrementally ℓ^2 stable*.

(b) $\hat{J}$ is a convex function of θ such that: (i) $\hat{J}(\theta, \tilde{z}) \geq J_{se}(\theta, x, \tilde{z})$ for any feasible x; (ii) $\hat{J}(\theta, \tilde{z})$ always has a finite minimum over θ; (iii) $\hat{J}(\theta, \tilde{z}) = 0$ when $\tilde{z}$ are generated by simulating the model represented by θ. In other words, the method perfectly recovers a "true" model if one exists in the model class.

A key tool we use is the Lagrangian relaxation, or S-Procedure, which is a method for approximating optimization problems with "difficult" nonconvex constraints by "easier" unconstrained problems with multipliers. Versions of this method have been used extensively in robust control and combinatorial optimization [25], [26]. In the context of this paper the idea is to approximate problems of the form

$$J^\star = \min_{x, \theta} J(\theta, x), \text{ subject to } h(\theta, x) = 0 \quad (4)$$

where $h(\theta, x)$ is a vector of constraint functions. In the context of simulation-error minimization, $J(\theta, x)$ represents the measure of model fit, e.g. (3), and $h(\theta, x)$ represents the dynamical constraints, e.g. (1).

Suppose that J is a convex function of θ and h is an affine function of θ, but both are nonconvex functions of θ and x jointly. In our context, Lagrangian relaxation refers to the approximation of the above problem by the following:

$$\hat{J}_\lambda^\star = \min_\theta \hat{J}_\lambda(\theta), \quad \hat{J}_\lambda(\theta) := \sup_x \{J(\theta, x) + \lambda(x)' h(\theta, x)\} \quad (5)$$

where λ is some vector function analogous to a Lagrange multiplier in nonlinear programming. The function $\hat{J}_\lambda(\theta)$ is clearly convex, since it is the supremum of a family of convex functions [27]. It is also an upper bound for $J(\theta, x)$ for any feasible x: if $h(\theta, x) = 0$ in (5), then $\hat{J}_\lambda(\theta) = J(\theta, x)$, so the supremum over x must be no smaller.

It is apparent that if the Lagrangian relaxation (5) is applied using two different constraint functions that nevertheless have the same feasible set: $h_1(\theta, x) = 0 \Leftrightarrow h_2(\theta, x) = 0$, but which differ for other values of θ, x, then the resulting values of $\hat{J}_\lambda(\theta)$ can differ. Consequently, performance of the Lagrangian relaxation may be improved by allowing a *redundant* parametrization of dynamical systems.

We will frequently deal with convex functions of the form $f(c, P) = c'P^{-1}c$, where c and P are decision variables, c is a vector and P is a symmetric positive-definite matrix of appropriate size. These functions can be represented by a slack variable and an LMI using the Schur complement [27]. Concave functions $g(c, P) = -c'P^{-1}c$ with $P > 0$ obey the following upper bound:

$$-c'P^{-1}c \leq b'Pb - 2b'c \tag{6}$$

where the right-hand side is a convex function of c, P for any fixed b. Inequality (6) follows directly from the expansion $b'Pb - 2b'c + c'P^{-1}c = |c - Pb|^2_{P^{-1}} \geq 0$. From this expansion it is also clear that the bound (6) is tight if $c = Pb$.

III. Convex Parameterization of Stable Models

The models we propose are constructed from three linearly parametrized vector functions:

$$e_\theta(x) = \sum_{i=0}^{q} \theta_i e_i(x), \qquad f_\theta(x, u) = \sum_{i=0}^{q} \theta_i f_i(x, u),$$
$$g_\theta(x, u) = \sum_{i=0}^{q} \theta_i g_i(x, u), \tag{7}$$

where $\theta \in \mathbb{R}^q$ and each basis function $e_i : \mathbb{R}^n \to \mathbb{R}^n, f_i : \mathbb{R}^n \times \mathbb{R}^m \to \mathbb{R}^n, g_i : \mathbb{R}^n \times \mathbb{R}^m \to \mathbb{R}^p$ is continuously differentiable in its arguments. These functions are associated with a parametrization of implicit state space dynamic models

$$e_\theta(x(t+1)) = f_\theta(x(t), u(t)), \tag{8}$$
$$y(t) = g_\theta(x(t), u(t)). \tag{9}$$

A particular choice of state dimension and bases for e, f, g defines a parametrization of models of the form (1), (2) given by $a(x, u, \theta) = e_\theta^{-1}(f_\theta(x, u))$ and $g(x, u, \theta) = g_\theta(x, u)$, as long as it can be guaranteed that $e_\theta(x)$ is a bijection for all θ.

We denote by $\mathcal{E}, \mathcal{F}, \mathcal{G}$ the span of the functions (7) over $\theta \in \mathbb{R}^q$. We assume that $\mathcal{E}, \mathcal{F}, \mathcal{G}$ contain all linear functions $x \mapsto Ex, (x, u) \mapsto Ax + Bu, (x, u) \mapsto Cx + Du$, respectively, where A, B, C, D, E are matrices of appropriate dimensions, and that $\mathcal{F}$ is closed under multiplication by a non-singular $n \times n$ matrix, i.e. if $f(\cdot, \cdot) \in \mathcal{F}$ then $Mf(\cdot, \cdot) \in F$ for all $M \in \mathbb{R}^{n \times n}$ with $\det(M) \neq 0$. Simple choices for $\mathcal{E}, \mathcal{F}, \mathcal{G}$ satisfying these assumptions are vectors of multivariate polynomials or trigonometric polynomials up to some fixed degree, parametrized by their coefficients. For the sake of brevity, we will usually drop the subscript θ from the functions e, f and g and speak of searching directly for e, f and g.

To ensure model stability, we propose the following constraint, requiring an auxiliary matrix variable $P = P' > 0$:

$$|f(x_1, u) - f(x_2, u)|^2_{(P^{-1})} - 2(x_1 - x_2)'(e(x_1) - e(x_2)) + |x_1 - x_2|^2_{P+\mu I} + |g(x_1, u) - g(x_2, u)|^2 \leq 0 \tag{10}$$

for all $x_1, x_2 \in \mathbb{R}^n, u \in \mathbb{R}^m$ for some fixed $\mu > 0$. Note that this constraint is convex with respect to e, f, g, P and hence convex with respect to a parameter vector θ when e, f, g, P are affinely parametrized.

Definition 2: For a given choice of state dimension and basis functions e_i, f_i, g_i, let Θ be the set of $\theta \in \mathbb{R}^q$ for which the model (8), (9) satisfies the inequality (10) for some $P = P' > 0$. Additionally, let $\hat{\mathcal{M}}$ be the model class $(e_\theta^{-1}(f_\theta(x, u)), g_\theta(x, u), \Theta)$.

Model well-posedness, i.e. existence of a unique solution, is guaranteed by the following theorem:

Theorem 1: Given a model (8), (9), if (10) holds for all x_1, x_2, u, then e is a bijection $\mathbb{R}^n \leftrightarrow \mathbb{R}^n$.

Proof. We need to show that for every $z \in \mathbb{R}^n$, the equation $e(x) = z$ has a unique solution. Firstly, note that (10) implies $2(x_1 - x_2)'(e(x_1) - e(x_2)) \geq \mu|x_1 - x_2|^2$. It is known (see, e.g., [28, Th. 18.15]) that if $T : \mathbb{R}^n \to \mathbb{R}^n$ is a nonlinear operator satisfying $(x_1 - x_2)'(T(x_1) - T(x_2)) \geq c|x_1 - x_2|^2$ for all $x_1, x_2 \in \mathbb{R}^n$ with $c > 0$, then the nonlinear equation $T(x) = 0$ has a unique solution. Take $T(x) = e(x) - z$ and $c = \mu/2$ and the theorem is proved. ∎

For simulation of the model, a solution to the equation $e(x) = z$ can be found by application of the Newton method. Convergence can be guaranteed by guarding with the ellipsoid method [24, Sec 2.2.2].

Theorem 2: For any model of the form (8), (9), the condition (10) implies global incremental ℓ^2 stability.

Proof. Consider the incremental Lyapunov function candidate $V(x_1, x_2) = |e(x_1) - e(x_2)|^2_{(P^{-1})}$ with $P = P' > 0$. The inequality

$$-|e(x_1) - e(x_2)|^2_{(P^{-1})} + |f(x_1, u) - f(x_2, u)|^2_{(P^{-1})} + \mu|x_1 - x_2|^2 + |g(x_1, u) - g(x_2, u)|^2 \leq 0 \tag{11}$$

for all $x_1, x_2 \in \mathbb{R}^n, u \in \mathbb{R}^m$ is a sufficient condition for global incremental ℓ^2 stability. Now apply the bound (6), with $c = e(x_1) - e(x_2)$ and $b = x_1 - x_2$, giving

$$-|e(x_1) - e(x_2)|^2_{(P^{-1})} \leq |x_1 - x_2|^2_P - 2(x_1 - x_2)'(e(x_1) - e(x_2)).$$

and substitute the right-hand-side of this inequality for $-|e(x_1) - e(x_2)|^2_{(P^{-1})}$ in (11). It is then clear that (10) implies (11) and hence stability of the model. ∎

Remark 1: It is straightforward to impose other behavioral constraints, e.g. passivity, by the same convexification procedure and appropriate choice of supply rate.

Definition 3: A system $x(t+1) = a(x(t), u(t)), y(t) = g(x(t), u(t))$, defined by functions $a : \mathbb{R}^n \times \mathbb{R}^m \mapsto \mathbb{R}^n$ and $g : \mathbb{R}^n \times \mathbb{R}^m \mapsto \mathbb{R}^p$ is called quadratically incrementally ℓ^2 stable, or more briefly: quadratically stable, if there exists an $M \in \mathbb{R}^{n \times n}$ with $M = M' > 0$ such that the condition

$$|a(x_1, u) - a(x_2, u)|^2_M - |x_1 - x_2|^2_M \leq -|g(x_1, u) - g(x_2, u)|^2, \tag{12}$$

holds for all $x_1, x_2 \in \mathbb{R}^n$ and $u \in \mathbb{R}^m$. I.e. incremental stability can be verified by a quadratic Lyapunov function.

Theorem 3: The model set $\hat{\mathcal{M}}$ includes all quadratically incrementally ℓ^2 stable models of the form (1), (2) for which $a(\cdot, \cdot) \in \mathcal{F}$ and $g(\cdot, \cdot) \in \mathcal{G}$.

Proof. Suppose (12) holds for some $M = M' > 0$. Choose $e(x) = Mx$, $f(x,u) = Ma(x,u)$, and $P = M$, i.e. $c = Pb$ in (6), and leave $g(x,u)$ as is. Clearly this model is equivalent to (1), (2), and by assumption that $a(\cdot,\cdot) \in \mathcal{F}$ and that $\mathcal{F}$ is closed under multiplication by M, $f(\cdot,\cdot) \in \mathcal{F}$. Direct substitution into (10) reveals that if (12) holds then so does (10). ∎

Remark 2: All stable linear systems are automatically stable in the above sense, as are models with stable linear dynamics but nonlinear input and/or output maps, e.g. Wiener and Hammerstein models, as long as $g(x,u)$ is globally Lipschitz in x.

IV. Convex Upper Bounds for Simulation Error

Let us suppose there is method for generating a vector sequence $\tilde{x} \in \ell_T^n$ from available data $\tilde{z}$. For our main theoretical results $\tilde{x}$ can be completely arbitrary, but it will be shown below in Theorem 4 that best results are obtained if $\tilde{x}$ approximates a good state sequence for the model, hence we refer to $\tilde{x}$ as the *surrogate state* sequence. Typical methods for generating $\tilde{x}$ include subspace identification [29], using a history of outputs as in nonlinear ARX [10], and – in the context of model reduction – proper orthogonal decomposition [2]. With some abuse of notation, from now on we consider the available data set to include $\tilde{x}$, i.e. $\tilde{z} := \{\tilde{u}, \tilde{y}, \tilde{x}\}$.

Given a surrogate state sequence, we will frequently refer to the following *equation error* functions:

$$\epsilon(\theta, \tilde{z}, t) = e_\theta(\tilde{x}(t+1)) - f_\theta(\tilde{x}(t), \tilde{u}(t)), \tag{13}$$

$$\eta(\theta, \tilde{z}, t) = \tilde{y}(t) - g_\theta(\tilde{x}(t), \tilde{u}(t)), \tag{14}$$

although for brevity of notation we suppress the arguments $(\theta, \tilde{z}, t)$ and simply denote dependence on time by subscripts: ϵ_t, η_t, and similarly for other signals e.g. $x_t = x(t)$.

Equation error ϵ_t is the one-step prediction error in a coordinate system generated by the bijection $e(\cdot)$. A straightforward way to search for a model is to minimize the equation error in a least-squares sense: $J_{EE}(\theta, \tilde{z}) = \sum_{t=0}^{T} |\epsilon_t|^2 + |\eta_t|^2$, which is clearly convex if models are affinely parameterized. E.g. [16] suggests identifying linear models in precisely this way. However, while this is often effective there is in general no direct relationship between J_{EE} and the (open-loop) simulation error, which results from recursion as described in Section II.

Application of the Lagrangian relaxation (5) suggests the following convex upper bound for simulation error, which is valid for any choice of n-dimensional vector function $\lambda(x, \tilde{z}, t)$:

$$\sup_{x \in l_T^n} \left\{ \sum_{t=0}^{T} |g(x_t, \tilde{u}_t) - \tilde{y}_t|^2 + \sum_{t=0}^{T-1} \lambda'\big(e(x_{t+1}) - f(x_t, \tilde{u}_t)\big) \right\}.$$

Our first proposal for a convex upper bound on simulation error comes from choosing $\lambda(x, \tilde{z}, t) = -2(x_{t+1} - \tilde{x}_{t+1})$:

$$\begin{aligned} \hat{J}_L(\theta, \tilde{z}) := \sup_{x \in l_T^n} \Bigg\{ & \sum_{t=0}^{T} |g(x_t, \tilde{u}_t) - \tilde{y}_t|^2 \\ & - \sum_{t=0}^{T-1} 2(x_{t+1} - \tilde{x}_{t+1})'(e(x_{t+1}) - f(x_t, \tilde{u}_t)) \Bigg\}, \\ & \text{subject to } x(0) = \tilde{x}(0). \end{aligned} \tag{15}$$

With this upper bound and the model set from Definition 2, we construct the following convex optimization problem.

Problem 1: For a model class (a, g, Θ) as in Definition 2 and data set $\tilde{z}$, solve $\hat{J}_L^\star(\tilde{z}) = \min_{\theta \in \Theta} \hat{J}_L(\theta, \tilde{z})$.

A challenge in solving Problem 1 is that it involves computing the supremum of a function of a large number of variables, the entire sequence x. An upper bound can be constructed by interchanging the order of supremum and summation, and introducing a positive-semidefinite storage function $V(x,t)$:

$$\begin{aligned} J(\theta, \tilde{z}) \leq \sum_{t=0}^{T} \sup_{x_t, x_{t+1}} \Big\{ & |g(x_t, \tilde{u}_t) - \tilde{y}_t|^2 \\ & + \lambda(x, \tilde{z}, t)'(e(x_{t+1}) - f(x_t, \tilde{u}_t)) \\ & + V(x_{t+1}, t+1) - V(x_t, t) \Big\}. \end{aligned} \tag{16}$$

Again, the quality of the bound will depend on the choice of functions V and λ.

It will be shown below that a particular choice of V and λ leads to the upper bound $\hat{J}_V(\theta, \tilde{z}) := \sum_{t=0}^{T} \mathcal{E}(\theta, \tilde{z}, t)$, where

$$\begin{aligned} \mathcal{E}(\theta, \tilde{z}, t) := \sup_{x \in \mathbb{R}^n} \{ & |f(x, \tilde{u}_t) - f(\tilde{x}_t, \tilde{u}_t) - \epsilon_t|^2_{(P^{-1})} + |x - \tilde{x}_t|^2_P \\ & - 2(x - \tilde{x}_t)'(e(x) - e(\tilde{x}_t)) + |g(x, \tilde{u}_t) - \tilde{y}_t|^2 \} \end{aligned} \tag{17}$$

and ϵ_t is given by (13). We define the following optimization problem, which is convex in θ.

Problem 2: For a model class (a, g, Θ) as in Definition 2 and data set $\tilde{z}$, solve $\hat{J}_V^\star(\tilde{z}) = \min_{\theta \in \Theta} \hat{J}_V(\theta, \tilde{z})$.

This upper bound is the *robust identification error* (RIE), first introduced in [21]. The following definition characterizes data sets for which "the true system is in the model class" [8].

Definition 4: Given a model class $\hat{\mathcal{M}} = (a, g, \Theta)$ and a data set $\tilde{z}$ we say that $\tilde{z}$ is in the *range* of $\hat{\mathcal{M}}$ if there exists a model in $\hat{\mathcal{M}}$ satisfying $\epsilon(t) = \eta(t) = 0$ for all $t = 0, 1, ...T$. That is, simulating this model with the observed input and initial conditions exactly reproduces the observed state and output.

We are now ready to state the main theoretical result regarding Problems 1 and 2.

Theorem 4: Given a model class $\hat{\mathcal{M}}$ and a data set $\tilde{z}$, Let $J_{se}^\star(\tilde{z})$ be the minimal simulation error over the model class $\hat{\mathcal{M}}$, then for any data set $\tilde{z}$ and any model parameters θ in $\hat{\mathcal{M}}$, $J_{se}(\theta, \tilde{z}) \leq \hat{J}_L(\theta, \tilde{z}) \leq \hat{J}_V(\theta, \tilde{z})$ and $\min_{\theta \in \hat{\mathcal{M}}} \hat{J}_V(\theta, \tilde{z}) < \infty$. Furthermore, if the data set $\tilde{z}$ is in the range of $\hat{\mathcal{M}}$, then $J_{se}^\star(\tilde{z}) = \hat{J}_L^\star(\tilde{z}) = \hat{J}_V^\star(\tilde{z}) = 0$.

Proof. The statement that $J_{se}(\theta, \tilde{z}) \leq \hat{J}_L(\theta, \tilde{z})$ for all θ follows from the basic properties of the Lagrangian relaxation.

To analyze $\hat{J}_V(\theta, \tilde{z})$ we first show that it corresponds to (16) with a particular choice of storage function and multiplier. Indeed, let $V(x,t) = 2(x - \tilde{x}_t)'(e(x) - e(\tilde{x}_t)) - |x - \tilde{x}_t|_P$ and $\lambda_t = -2(x_{t+1} - \tilde{x}_{t+1})$, then (16) reduces to:

$$\begin{aligned} \sup_{x_t, x_{t+1}} \{ & |g(x_t, \tilde{u}_t) - \tilde{y}_t|^2 - |x_{t+1} - \tilde{x}_{t+1}|_P \\ & + 2(x_{t+1} - \tilde{x}_{t+1})'(f(x_t, \tilde{u}_t) - f(\tilde{x}_t, \tilde{u}_t) - \epsilon_t) \\ & - 2(x_t - \tilde{x}_t)'(e(x_t) - e(\tilde{x}_t)) + |x_t - \tilde{x}_t|_P \}. \end{aligned} \tag{18}$$

With this formulation, x_{t+1} only appears in quadratic or linear terms, so the partial maximum can be computed exactly: $|f(x_t, \tilde{u}_t) - f(\tilde{x}_t, \tilde{u}_t) - \epsilon_t|^2_{(P^{-1})} = \max_{x_{t+1}} \{ -|x_{t+1} - \tilde{x}_{t+1}|^2_P$

$+ 2(x_{t+1} - \tilde{x}_{t+1})'(f(x_t, \tilde{u}_t) - f(\tilde{x}_t, \tilde{u}_t) - \epsilon_t)\}$ giving the following bound in place of (18): $\sup_{x(t)}\{|g(x_t, \tilde{u}_t) - \tilde{y}_t|^2 + |f(x_t, \tilde{u}_t) - f(\tilde{x}_t, \tilde{u}_t) - \epsilon_t|^2_{(P^{-1})} - 2(x_t - \tilde{x}_t)'(e(x_t) - e(\tilde{x}_t)) + |x_t - \tilde{x}_t|_P\}$ which is $\mathcal{E}(\theta, \tilde{z}, t)$ in (17).

To prove the statement $\hat{J}_L(\theta, \tilde{z}) \leq \hat{J}_V(\theta, \tilde{z})$, we first apply (16) with the choice of V and λ given above, to give

$$\sup_x \left\{ \sum_{t=0}^{T} \big(V(x_{t+1}, t+1) - V(x_t, t) + |g(x_t, \tilde{u}_t) - \tilde{y}_t|^2 + \lambda_t'(e(x_{t+1}) - f(x_t, \tilde{u}_t)) \big) \right\} \leq \hat{J}_V(\theta, \tilde{z}). \quad (19)$$

All storage function terms cancel in the summation except the first and the last. By construction[2], with $x(0) = \tilde{x}(0)$ we have $V(x(0), 0) = 0$, and since V is positive-definite $V(x(T), T) \geq 0$, the left-hand side of (19) is an upper bound for $\sup_x \left\{ \sum_{t=0}^{T} \left(|g(x_t, \tilde{u}_t) - \tilde{y}_t|^2 + \lambda_t'(e(x_{t+1}) - f(x_t, \tilde{u}_t)) \right) \right\}$ which equals $\hat{J}_L$ plus an additional multiplier term. This term clearly does not decrease the supremum, since the choice $x_{t+1} = e^{-1}(f(x_t, \tilde{u}_t))$ would make it zero for any choice of x_t. This proves that $\hat{J}_L(\theta, \tilde{z}) \leq \hat{J}_V(\theta, \tilde{z})$.

To prove that the minimal value of $\hat{J}_V(\theta, \tilde{z})$ is finite, note that one can choose a model (7) with $f(x, u) = 0, g(x, u) = 0$. In this case, the bound (17) reduces to $\mathcal{E}(\theta, \tilde{z}, t) = \sup_{x \in \mathbb{R}^n}\{|\tilde{y}_t|^2 + |\epsilon_t|^2_{(P^{-1})} + |x - \tilde{x}_t|^2_P - 2(x - \tilde{x}_t)'(e(x) - e(\tilde{x}_t))\}$ and the stability constraint (10) ensures that $|x - \tilde{x}_t|^2_P - 2(x - \tilde{x}_t)'(e(x) - e(\tilde{x}_t)) \leq -\mu|x - \tilde{x}_t|^2$ for all x, t. Since $\tilde{y}_t$ and ϵ_t are not functions of x this proves the supremum is finite.

Similarly, if $\tilde{z}$ is in the range of $\hat{\mathcal{M}}$, then there exists a model $\bar{\theta} \in \hat{\mathcal{M}}$ for which $\epsilon(t) = 0, \eta(t) = 0$ for all t. Let $\bar{e}, \bar{f}, \bar{g}$ be the model functions given by $\bar{\theta}$, then (17) is:

$$\mathcal{E}(\bar{\theta}, \tilde{z}) = \sup_{x \in \mathbb{R}^n} \{|\bar{g}(x, \tilde{u}_t) - \bar{g}(\tilde{x}_t, \tilde{u}_t)|^2 + |x - \tilde{x}_t|^2_P + |\bar{f}(x, \tilde{u}_t) - \bar{f}(\tilde{x}_t, \tilde{u}_t)|^2_{(P^{-1})} - 2(x - \tilde{x}_t)'(\bar{e}(x) - \bar{e}(\tilde{x}_t))\}$$

and by (10) this can be bounded by $\mathcal{E}(\bar{\theta}, \tilde{z}, t) \leq \sup_{x \in \mathbb{R}^n}\{-\mu|x - \tilde{x}_t|^2\} = 0$. Therefore $\hat{J}_V(\tilde{z}, \bar{\theta}) = 0$. Now, since $0 \leq J_{se}(\theta, \tilde{z}) \leq \hat{J}_L(\theta, \tilde{z}) \leq \hat{J}_V(\theta, \tilde{z})$ for all θ, this model must give the minimal value of $\hat{J}_V(\tilde{z}, \bar{\theta})$, and furthermore $J(\tilde{z}, \bar{\theta}) = 0$ so a perfect model is recovered. Note that we do not require existence of a *unique* perfect model. ■

V. Contracting Models and Linearized Simulation Error

The upper bounds in the previous section, while convex, may be expensive to compute when the dependence of e, f, and g on x is nonlinear. Contraction analysis [30] provides an alternative set of tools for studying model behaviour, based on infinitesimal displacements of the system rather than pairs of solutions. The principal advantage in the context of this paper is that it results in significantly easier optimization problems involving suprema of *quadratic* functions, which can be computed explicitly or represented using LMIs.

[2]Alternatively, when working without a known initial condition, one can remove the constraint $x(0) = \tilde{x}(0)$ and add $-V(x_0, \tilde{x}_0) = -2(x_0 - \tilde{x}_0)'(e(x_0) - e(\tilde{x}_0)) + |x_0 - \tilde{x}_0|_P$ to terms inside the brackets in (15).

A. Contraction Condition for Incremental ℓ^2 Stability

In this section we derive contraction conditions for incremental ℓ^2 stability. Contraction analysis is based on the study of an extended system consisting of (8), (9) as well as the *differential* dynamics

$$E(x(t+1))\Delta(t+1) = F(x(t), u(t))\Delta(t), \quad (20)$$
$$\Delta_y(t) = G(x(t), u(t))\Delta(t), \quad (21)$$

i.e. the dynamics (8), (9) linearized with respect to initial conditions, where $E(x) = \frac{\partial}{\partial x}e(x)$, $F(x, u) = \frac{\partial}{\partial x}f(x, u)$, and $G(x, u) = \frac{\partial}{\partial x}g(x, u)$.

For our purposes, a differential storage function can be any positive-semidefinite function of x and Δ, i.e. $V : \mathbb{R}^n \times \mathbb{R}^n \to \mathbb{R}$ such that $V(x, \Delta) \geq 0$ for all x and Δ. If a differential storage function can be found that verifies the "differential" dissipation inequality

$$V(x(t+1), \Delta(t+1)) - V(x(t), \Delta(t)) \leq - |G(x(t), u(t))\Delta_t|^2 - \mu|\Delta(t)|^2 \quad (22)$$

for solutions of (8), (20), then $V(x(0), \Delta(0))$ is an upper bound for $\sum_{t=0}^{T} |G(x(t), u(t))\Delta_t|^2 + \mu\sum_{t=0}^{T} |\Delta(t)|^2$ for all T, for all solutions of the extended system (8), (9), (20), (21).

We propose to use the Riemannian metric $V(x, \Delta) = |E(x)\Delta|^2_{(P^{-1})}$ with $P = P' > 0$ as a differential storage function, with which (22) becomes

$$|F(x, u)\Delta|^2_{(P^{-1})} - |E(x)\Delta|^2_{(P^{-1})} + |G(x, u)\Delta|^2 \leq -\mu|\Delta|^2. \quad (23)$$

Just as (11) is implied by (10), (23) is implied by

$$F(x, u)'P^{-1}F(x, u) + P - E(x) - E(x)' + G(x, u)'G(x, u) \leq -\mu I. \quad (24)$$

Furthermore, unlike (23), condition (24) is jointly convex in e, f, g and P, i.e. convex in the parameter vector θ. We now give an alternative model class to $\hat{\mathcal{M}}$ in Section III.

Definition 5: For a given choice of state dimension and basis functions e_i, f_i, g_i, we define Θ^0 as the set of $\theta \in \mathbb{R}^q$ for which the model (8), (9) satisfies the inequality (24) for some $P = P' > 0$. Additionally, we define the model class $\hat{\mathcal{M}}^0 := (e_\theta^{-1}(f_\theta(x, u)), g_\theta(x, u), \Theta^0)$.

The main theoretical result for this model class is the following, which is analogous to Theorems 1, 2, and 3.

Theorem 5: All models in $\hat{\mathcal{M}}^0$ are well-posed and incrementally ℓ^2 stable. Addtionally, the set $\hat{\mathcal{M}}^0$ contains all models of the form (1), (2) for which (22) can be verified with a constant quadratic metric $V(x, \Delta) = \Delta' M \Delta$ for some $M = M' > 0$ and $a(x, u) \in \mathcal{F}$.

Proof. To prove well-posedness, consider any pair $x_1, x_2 \in \mathbb{R}^n$ and define $x_\rho = (1 - \rho)x_1 + \rho x_2$ with $\rho \in [0, 1]$. Every term on the left-hand side of (24) is positive semidefinite except $-E(x) - E(x)'$, so $E(x) + E(x)' \geq \mu I$. Now, $(x_2 - x_1)'(e(x_2) - e(x_1)) = \int_0^1 (x_2 - x_1)'E(x_\rho)(x_2 - x_1)d\rho \geq \frac{\mu}{2}|x_2 - x_1|^2$, so well-posedness is guaranteed by the same argument as in the proof of Theorem 1.

To prove incremental stability, consider any pair of initial conditions $x_1(0)$ and $x_2(0)$, and consider the family of solutions to (8) x_ρ, parameterized by $\rho \in [0, 1]$, having initial

conditions $x_\rho(0) = (1-\rho)x_1(0) + \rho x_2(0)$ and any fixed input sequence $\bar{u}(t)$. Note that for $t > 0$, $x_\rho(t)$ does not necessarily linearly interpolate $x_1(t)$ and $x_2(t)$.

Now, define $\Delta_\rho(t) := \frac{\partial}{\partial \rho} x_\rho(t)$. For any $T > 0$ we have $\sum_{t=0}^{T} |y_2(t) - y_1(t)|^2 = \sum_{t=0}^{T} \left| \int_0^1 G(x_\rho(t))\Delta_\rho(t) d\rho \right|^2 \leq \int_0^1 \sum_{t=0}^{T} |G(x_\rho(t))\Delta_\rho(t)|^2 \, d\rho$, by the Cauchy-Schwarz inequality and interchanging order of summation and integration.

Summing (23) repeatedly with $x = x_\rho(t)$, $u = \bar{u}(t)$ and $\Delta = \Delta_\rho(t)$ gives $\sum_{t=0}^{T} |y_2(t) - y_1(t)|^2 \leq \int_0^1 |E(x_\rho(0))\Delta_\rho(0)|^2_{(P^{-1})} - |E(x_\rho(T+1))\Delta_\rho(T+1)|^2_{(P^{-1})} \, d\rho$
$\leq \int_0^1 |E(x_\rho(0))(x_2(0) - x_1(0))|^2_{(P^{-1})} \, d\rho$, as $\Delta_\rho(0) = x_2(0) - x_1(0)$. As this bound holds uniformly for all T we arrive at the desired result. Exactly the same reasoning can be used to bound $\sum_{t=0}^{T} |x_2(t) - x_1(t)|^2$.

The proof that all models with a constant quadratic metric are in $\hat{\mathcal{M}}^0$ is similar to the proof of Theorem 3, and follows directly from the choice $e(x) = Mx, f(x,u) = Ma(x,u)$. ■

Condition (24) is, after taking the Schur complement of $F(x,u)'P^{-1}F(x,u)$, a *point-wise* linear matrix inequality, i.e. for every pair (x,u) there is an LMI. For polynomial models, this can be enforced using a sum-of-squares constraint [31].

B. Linearized Simulation Error and its Convex Bounds

We now introduce an approximation to simulation error that is easier to optimize. Consider a "perturbed" version of the model equations, motivated by (13), (14):

$$e(x_\rho(t+1)) = f(x_\rho(t), u(t)) + \rho\epsilon(t), \quad (25)$$
$$y_\rho(t) = g(x_\rho(t), u(t)) + \rho\eta(t), \quad (26)$$

where $\rho \in [0,1]$ and $\epsilon(t), \eta(t)$ are the equation errors.

By construction we have $y_0(t) = y(t)$, the simulated output, and $y_1(t) = \tilde{y}(t)$, the measured output. We define the linearized simulation error as

$$J^0(\theta, \tilde{z}) := \lim_{\rho \to 1} \frac{1}{(1-\rho)^2} \sum_{t=0}^{N} |\tilde{y}(t) - y_\rho(t)|^2 \quad (27)$$

to quantify local sensitivity of model equations with respect to equation errors. Applying L'Hospital's rule twice to (27) we arrive at the alternative expression $J^0(\theta, \tilde{z}) = \sum_{t=0}^{N} |G(\tilde{x}(t), \tilde{u}(t))\tilde{\Delta}(t) + \eta(t)|^2$, where $\tilde{\Delta}(t) = \left.\frac{\partial x_\rho(t)}{\partial \rho}\right|_{\rho=1}$ satisfies $\tilde{\Delta}(0) = 0$ and

$$E(\tilde{x}(t+1))\tilde{\Delta}(t+1) = F(\tilde{x}(t), \tilde{u}(t))\tilde{\Delta}(t) + \epsilon(t). \quad (28)$$

Roughly speaking, when the simulation error is small, i.e. the model fit is good, the simulation error and the linearized simulation error will be very close, and $\tilde{\Delta}(t) \approx \tilde{x}(t) - x(t)$. Note that if the model equations are affine in x, then the simulation error and the linearized simulation error are identical.

A family of convex upper bounds for the linearized simulation error can be constructed in a similar way to those in Section IV, so we skip some details in the construction.

The analogue of Problem 1 can be given in a simple form by noting that (28) can be written in "lifted" representation $H(\theta, \tilde{z})\vec{\Delta} = \vec{\epsilon}(\theta, \tilde{z})$ with the stacked vectors $\vec{\Delta} := [\tilde{\Delta}(1)', \tilde{\Delta}(2)', ..., \tilde{\Delta}(T)']'$ and $\vec{\epsilon}(\theta, \tilde{z}) := [\epsilon(0)', \epsilon(1)', ...\epsilon(T)']'$ and $H(\theta, \tilde{z})$ is a sparse matrix that is affine in θ and is straightforward to construct from (28). The analogue of (15) can then be written as

$$\hat{J}^0_L(\theta, \tilde{z}) := \sup_{\vec{\Delta} \in l_T^{nT}} \{|\vec{G}(\tilde{z})\vec{\Delta} + \vec{\eta}(\tilde{z})|^2 - 2\vec{\Delta}'(H(\theta, \tilde{z})\vec{\Delta} - \vec{\epsilon}(\tilde{z}))\}$$

where $\vec{G}(\tilde{z})$ is a block-diagonal matrix with elements $G(\tilde{z}(t))$ and $\vec{\eta}$ is a stacked vector of output equation errors $\eta(t)$. Note that both H and $\vec{\epsilon}$ are affine in θ, so the functional $\hat{J}^0_L$ is convex in θ. Note also that the supremum is taken of a quadratic function of $\vec{\Delta}$, which will be shown below to be concave, and can therefore be computed explicitly.

We can also construct an analogue of Problem 2 for linearized simulation error. Define the pointwise functions

$$\mathcal{E}^0(\tilde{z}, t) := \sup_{\Delta \in \mathbb{R}^n} \{|F(\tilde{x}, \tilde{u})\Delta + \epsilon|^2_{(P^{-1})} + |\Delta|^2_P - 2\Delta' E(\tilde{x})\Delta + |G(\tilde{x}, \tilde{u})\Delta + \eta|^2\}, \quad (29)$$

and the following convex cost function: $\hat{J}^0_V(\theta, \tilde{z}) = \sum_{t=0}^{N} \mathcal{E}^0(\tilde{z}, t)$, which we call the *local robust identification error* (local RIE).

Theorem 6: For all $P = P' > 0$ and signal data $\tilde{z}$, for every model satisfying (24), $J^0(\theta, \tilde{z}) \leq \hat{J}^0_L(\theta, \tilde{z}) \leq \hat{J}^0_V(\theta, \tilde{z}) < \infty$. Furthermore, if the data set $\tilde{z}$ is in the range of $\hat{\mathcal{M}}$, then the optimal values satisfy $J^{0,\star}(\tilde{z}) = \hat{J}^{0,\star}_L(\tilde{z}) = \hat{J}^{0,\star}_V(\tilde{z}) = 0$.

Proof. First we show that the bound $\hat{J}^0_V$ is finite. Each term $\mathcal{E}^0(\tilde{z})$ is the supremum of quadratic functional of Δ, so each supremum is finite if and only if the corresponding matrix

$$F(\tilde{x}(t), \tilde{u}(t))'P^{-1}F(\tilde{x}(t), \tilde{u}(t)) + P - E(\tilde{x}(t)) - E(\tilde{x}(t))' + G(\tilde{x}(t), \tilde{u}(t))'G(\tilde{x}(t), \tilde{u}(t))$$

is negative semidefinite. In fact, the contraction condition (23) implies that the above matrix is is uniformly negative definite, which implies a finite supremum achieved by a unique Δ.

The remainder of the proof follows from the same arguments as the proof of Theorem 4, but with the differential storage function $V(\tilde{x}, \Delta) = \Delta'(E(\tilde{x}) + E(\tilde{x})' - P)\Delta$ and by substituting $\tilde{\Delta}_t$ for $\tilde{x}_t - x_t$, $E(\tilde{x}_t)\tilde{\Delta}_t$ for $e(\tilde{x}_t) - e(x_t)$ and $F(\tilde{x}_t)\tilde{\Delta}_t$ for $f(\tilde{x}_t, u_t) - f(x_t, u_t)$. ■

VI. Examples and Discussion

In this section we demonstrate the performance of the proposed methods on two real-world applications: reduced-order dynamic inversion of an op-amp and identification of a pneumatic actuator from experiments. Further examples illustrating the method can be found in [24].

We judge the performance of a model over T time samples using percent normalized simulation error: $J_{\text{perf}} = 100\sqrt{\sum_{t=0}^{T} |\tilde{y}(t) - y(t)|^2} / \sqrt{\sum_{t=0}^{T} |\tilde{y}(t) - \frac{1}{T+1}\sum_{s=0}^{T} \tilde{y}(s)|^2}$ where $y(t)$ is the result of simulating the model from observed initial conditions and with the known input. When minimizing the equation error, to ensure the dynamics are well-posed, we require $E(x) + E(x)' - 2\mu I$ to be positive semidefinite via a sums-of-squares constraint [31]. Software

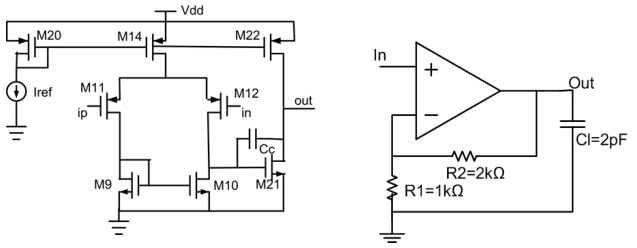

(a) Simplified schematic of opamp. (b) Opamp configuration.

Fig. 1. Simplified internal schematic of the opamp and feedback configuration used for data generation. Parasitics not shown.

TABLE I
PERFORMANCE (J_{PERF}) FOR OP-AMP DYNAMIC INVERSION USING LOCAL RIE WITH MODELS OF VARYING ORDER n AND DEGREE d.

$n \setminus d$	1	3	5
1	86.47	94.162	94.22
2	23.82	21.90	21.69
3	11.80	8.31	7.99

packages such as the Yalmip toolbox [32] can be used to transform polynomial optimization problems such as those in this paper to semidefinite programs. The examples in this paper took up to ten minutes to solve using the free SDP solver Sedumi [33], though preliminary work on specialised algorithms indicates substantial speed-ups are possible [34].

A. Low-Order Nonlinear Dynamic Inversion of an Op-Amp

Our first example is causal dynamic inversion of an op-amp using data from a detailed high-dimensional simulation model. Figure 1 presents the schematic of the op-amp at the transistor level, and the device in the closed loop simulation used for identification. For the task of dynamic inversion, we seek to recover v_{in} from v_{out}. To accomplish this we identify a dynamical system with v_{out} as its input and v_{in} as its output.

The identification data-set consists of two trials. Each trial is a simulation in response to v_{in} being the sum of five sinusoids. The frequencies were chosen randomly within the bandwidth of the amplifier (60 MHz). The trials have different biases and amplitudes, and excite saturation in the amplifier.

Model complexity is varied in terms of state and input dimension, and polynomial degree. Our input and surrogate state signals are in "autoregressive moving-average" form, parameterized by a positive integer n: $\tilde{y}(t) = v_{\text{in}}(t)$, $\tilde{u}(t) = \begin{bmatrix} v_{\text{out}}(t) & \dots & v_{\text{out}}(t+1-n) \end{bmatrix}'$, $\tilde{x}(t) = \begin{bmatrix} v_{\text{in}}(t) & \dots & v_{\text{in}}(t+1-n) \end{bmatrix}'$. We search over models with f being affine in the input and e, f of varying degree d in the state $x(t)$, with g fixed $g(x(t), u(t)) = x_1(t)$. For all such models, minimizing equation error J_{EE} resulted in an unstable model. Table I presents the results using RIE for varying model complexities. The minimizers of the local RIE dramatically improve in performance as the model complexity increases.

B. Identification of a Pneumatic Piston from Experiment

Our second example is the identification of a double-acting pneumatic cylinder under conditions simulating heavy load.

TABLE II
COMPARISON OF PERFORMANCE (J_{PERF}) FOR MODELS OF THE PNEUMATIC PISTON ON VALIDATION DATA.

Model Class	Equation Error	Local RIE
Linear	29.31	24.62
Polynomial ($\deg_x(e)$=1, $\deg_x(f)$=1)	14.34	22.07
Polynomial ($\deg_x(e)$=3, $\deg_x(f)$=1)	81.26	14.84
Polynomial ($\deg_x(e)$=3, $\deg_x(f)$=3)	Divergent	6.64

Each chamber is separated from a control valve by approximately 25 cm of tubing. A compressor supplies high pressure (between 0.6 and 0.8 MPa) which is then regulated to a 0.5 MPa supply pressure to the valves, which have internal pressure control mechanisms. Both ends of the piston are fixed to a bracket. The dynamics of the system include changes in supply pressure due to flow, filling the chambers, the internal dynamics of the valves and elastic bending of the rail to which the piston and bracket are fixed. The pressure is measured at the inlet of both chambers (denoted $p_1(t)$ and $p_2(t)$ resp.). A load cell between the bracket and piston measures the force exerted on the bracket ($f_o(t)$). We denote the commanded pressure signals as $v_1(t)$ and $v_2(t)$ for valve 1 and 2 resp. Finally, the supply pressure is measured near the inlet of the valves ($p_s(t)$). The input to the system is a scalar signal $f_d(t)$ which represents the desired force in Newtons.

The experiment consisted of the application of seven chirp signals of varying amplitudes being applied to the system. These chirps signals swept linearly in instantaneous frequency from approximately 3 to 40 Hz and peak desired force in the range 20 to 100 N. The identification procedures were trained on data sets with $20, 60, 100$ and 130 N being the peak absolute desired force. These models were then validated on trials with $40, 80$ and 120 N being the peak forces desired. The experiments were performed at a sample-rate of 1 kHz. The desired and realized force were taken to be the input and output signals and the internal pressures and output force were chosen as the surrogate states: $\tilde{u}(t) = \begin{bmatrix} f_d(t) & f_d(t-1) \end{bmatrix}'$, $\tilde{y}(t) = f_o(t), \tilde{x}(t) = \begin{bmatrix} p_1(t) & p_2(t) & p_s(t) & f_o(t) \end{bmatrix}'$. We fit several discrete time models of the form (8), (9). We examine polynomial models of increasing complexity, beginning with a linear model then input affine models with various degrees in x. In all cases the output map $g(\cdot, \cdot)$ is fixed to $g(x(t), u(t)) = x_4(t)$.

The computational complexity of minimizing the RIE scales with the number of data-points for nonlinear fits. As a result, we sub-select 10,000 data-points for optimizing the local RIE. To ensure a fair comparison in these cases, we fit an equation error model on both the same 10,000 data-points and the whole data-set and report the better of the two scores. Table II compares the results, in order of increasing model complexity. Figure 2 compares the observed $\tilde{y}(t)$ to the open loop simulation of the best performing model from the local RIE. Note that the two plots are responses to differently scaled versions of the same input, so nonlinearity is clearly apparent.

We observe that the local RIE under-performs equation error for the first polynomial model structure. However, with increasing model complexity the equation error minimization at first improves, then rapidly degrades in performance and becomes unstable. On the other hand, the performance of

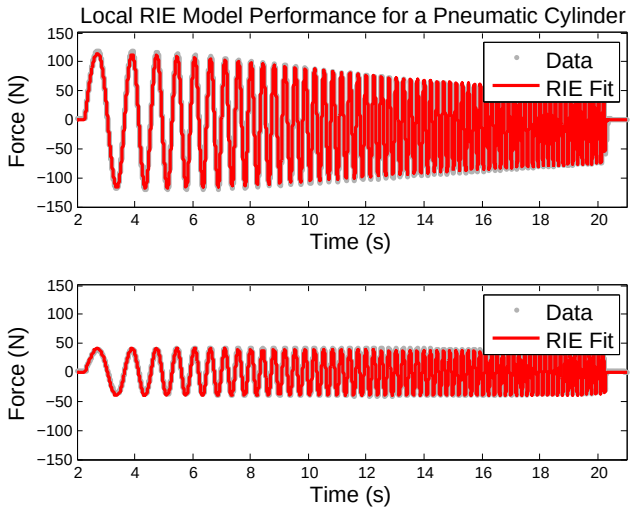

(a) Data and Simulation.

Fig. 2. Model performance using the local RIE: simulation (solid-red) compared against observed output (gray dots) for two different input amplitudes from nearly identical initial conditions. The response does not simply scale, indicating nonlinearity.

the local RIE minimization steadily improves with increasing model complexity. These observations are consistent with the op-amp dynamic inversion.

One interpretation of this performance is that the stability constraints and the "robust" fidelity bounds act as a form of regularization, and mitigate the usual harmful effects of over-fitting and reduce the need for structure selection or "pruning" of regressors in autoregressive nonlinear models [10]. Indeed, for linear systems we have observed that these constraints confer a bias towards "more stable" models [22], [34]. Quantifying this effect will be a topic of future research.

VII. Acknowledgments

We would like to thank Ryuma Niiyama for assistance with the pneumatic cylinder experiments, and Zhipeng Li, Yan Li and Vladimir Marko Stojanovic for their collaboration in providing the RF amplifier data set.

References

[1] A. C. Antoulas, *Approximation of large-scale dynamical systems*. Society for Industrial and Applied Mathematics, 2005, vol. 6.

[2] P. Holmes, J. Lumley, and G. Berkooz, *Turbulence, Coherent Structures, Dynamical Systems and Symmetry*. Cambridge University Press, 1998.

[3] S. Lall, J. E. Marsden, and S. Glavaski, "A subspace approach to balanced truncation for model reduction of nonlinear control systems," *International Journal of Robust and Nonlinear Control*, vol. 12, no. 6, pp. 519–535, 2002.

[4] D. J. Lucia, P. S. Beran, and W. A. Silva, "Reduced-order modeling: new approaches for computational physics," *Progress in Aerospace Sciences*, vol. 40, no. 1, pp. 51–117, 2004.

[5] J. Sjöberg, Q. Zhang, L. Ljung, A. Benveniste, B. Delyon, P.-Y. Glorennec, H. Hjalmarsson, and A. Juditsky, "Nonlinear black-box modeling in system identification: a unified overview," *Automatica*, vol. 31, no. 12, pp. 1691–1724, 1995.

[6] L. Ljung, *System Identification: Theory for the User*, 3rd ed. Englewood Cliffs, New Jersey, USA: Prentice Hall, 1999.

[7] R. K. Pearson, "Nonlinear empirical modeling techniques," *Computers & chemical engineering*, vol. 30, no. 10, pp. 1514–1528, 2006.

[8] L. Ljung, "Perspectives on system identification," *Annual Reviews in Control*, vol. 34, no. 1, pp. 1 – 12, 2010.

[9] T. Schön, A. Wills, and B. Ninness, "System identification of nonlinear state-space models," *Automatica*, vol. 47, no. 1, pp. 39–49, 2010.

[10] S. Billings and S. Fakhouri, "Identification of systems containing linear dynamic and static nonlinear elements," *Automatica*, vol. 18, no. 1, pp. 15–26, 1982.

[11] S. Billings, *Nonlinear System Identification: NARMAX Methods in the Time, Frequency, and Spatio-Temporal Domains*. Wiley, 2013.

[12] F. J. Doyle, R. K. Pearson, and B. A. Ogunnaike, *Identification and control using Volterra models*. Springer Verlag, 2002.

[13] P. A. Regalia and P. Stoica, "Stability of multivariable least-squares models," *IEEE Signal Processing Letters*, vol. 2, no. 10, pp. 195–196, 1995.

[14] J. M. Maciejowski, "Guaranteed stability with subspace methods," *Systems & Control Letters*, vol. 26, no. 2, pp. 153–156, 1995.

[15] T. Van Gestel, J. A. Suykens, P. Van Dooren, and B. De Moor, "Identification of stable models in subspace identification by using regularization," *IEEE Transactions on Automatic Control*, vol. 46, no. 9, pp. 1416–1420, 2001.

[16] S. L. Lacy and D. S. Bernstein, "Subspace identification with guaranteed stability using constrained optimization," *IEEE Transactions on Automatic Control*, vol. 48, no. 7, pp. 1259–1263, 2003.

[17] B. Besselink, N. van de Wouw, and H. Nijmeijer, "Model reduction for nonlinear systems with incremental gain or passivity properties," *Automatica*, vol. 49, no. 4, pp. 861–872, 2013.

[18] T. Söderström, "On the uniqueness of maximum likelihood identification," *Automatica*, vol. 11, no. 2, pp. 193–197, 1975.

[19] A. Megretski, "Convex optimization in robust identification of nonlinear feedback," in *Proceedings of the 47th IEEE Conference on Decision and Control*, Cancun, Mexico, Dec 9-11 2008, pp. 1370–1374.

[20] B. Bond, Z. Mahmood, Y. Li, R. Sredojevic, A. Megretski, V. Stojanovic, Y. Avniel, and L. Daniel, "Compact modeling of nonlinear analog circuits using system identification via semidefinite programming and incremental stability certification," *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, vol. 29, no. 8, pp. 1149 –1162, aug. 2010.

[21] M. M. Tobenkin, I. R. Manchester, J. Wang, A. Megretski, and R. Tedrake, "Convex optimization in identification of stable non-linear state space models," in *IEEE Conference on Decision and Control (CDC)*, 2010, pp. 7232–7237.

[22] I. R. Manchester, M. M. Tobenkin, and A. Megretski, "Stable nonlinear system identification: Convexity, model class, and consistency," in *IFAC Symposium on System Identification*, vol. 16, no. 1, 2012, pp. 328–333.

[23] M. M. Tobenkin, I. R. Manchester, and A. Megretski, "Stable nonlinear identification from noisy repeated experiments via convex optimization," in *American Control Conference (ACC)*, 2013, pp. 3936–3941.

[24] M. M. Tobenkin, "Robustness analysis for identification and control of nonlinear systems," Ph.D. dissertation, Massachusetts Institute of Technology, 2014.

[25] I. Pólik and T. Terlaky, "A survey of the S-lemma," *SIAM review*, vol. 49, no. 3, pp. 371–418, 2007.

[26] C. Lemaréchal, "Lagrangian relaxation," in *Computational Combinatorial Optimization*. Springer, 2001, pp. 112–156.

[27] S. Boyd and L. Vandenberghe, *Convex optimization*. Cambridge university press, 2009.

[28] A. Poznyak, *Advanced Mathematical Tools for Control Engineers: Volume 1: Deterministic Systems*. Elsevier, 2010.

[29] P. Van Overschee and B. De Moor, *Subspace identification for linear systems: theory, implementation, applications*. Kluwer academic publishers, 1996.

[30] W. Lohmiller and J.-J. E. Slotine, "On contraction analysis for non-linear systems," *Automatica*, vol. 34, no. 6, pp. 683–696, June 1998.

[31] P. A. Parrilo, "Semidefinite programming relaxations for semialgebraic problems," *Mathematical programming*, vol. 96, pp. 293–320, 2003.

[32] J. Lofberg, "Yalmip: A toolbox for modeling and optimization in matlab," in *Computer Aided Control Systems Design, 2004 IEEE International Symposium on*. IEEE, 2004, pp. 284–289.

[33] J. F. Sturm, "Using SeDuMi 1.02, a matlab toolbox for optimization over symmetric cones," *Optimization methods and software*, vol. 11, no. 1-4, pp. 625–653, 1999.

[34] J. Umenberger and I. R. Manchester, "Specialized algorithm for identification of stable linear systems using Lagrangian relaxation," in *Proc. 2016 American Control Conference*, Boston, MA, 2016.